

Traffic Service Position System No. 1B:

Capacity and Reliability Evaluation

By B. A. CRANE and D. S. SUK

(Manuscript received June 30, 1982)

This paper describes the prediction and evaluation of the call-processing capacity and reliability resulting from use of the 3B20 Duplex (3B20D) Processor in the Traffic Service Position System No. 1B (TSPS No. 1B). The call-processing capacity was predicted using a processor real-time model whose parameter values were determined by laboratory and test-site measurements. The system reliability was predicted using Markov modeling techniques. Performing an evaluation during TSPS No. 1B development provided a means for monitoring progress toward meeting the capacity and reliability objectives.

I. INTRODUCTION

One of the development objectives of the Traffic Service Position System No. 1B (TSPS No. 1B) was to improve the call-handling capacity of the TSPS No. 1 by replacing the Stored Program Control No. 1A (SPC 1A) with the Stored Program Control No. 1B (SPC 1B).¹ The SPC 1B consists of the 3B20 Duplex (3B20D) Processor together with the Peripheral System Interface (PSI) unit, which adapts the 3B20D to existing TSPS peripherals. The 3B20D is microprogrammed to execute the SPC 1A instructions, thus allowing TSPS call-processing software developed for the SPC 1A to be ported to the SPC 1B with minimal changes. This emulated TSPS software executes as a kernel process under the DMERT operating system. References 2 and 3 contain further details on the SPC 1B architecture.

1.1 Call-processing capacity analysis

The increased speed of the SPC 1B in executing the emulated SPC 1A instructions provides the increase in call-processing capacity. The

initial objective for capacity increase established for TSPS No. 1B was that the SPC 1B call-processing capacity should be at least 160 percent that of the SPC 1A.

Early in TSPS No. 1B development, a capacity prediction and evaluation plan was established for monitoring the progress in meeting the capacity improvement objective. This plan involved formulating a mathematical model of the SPC 1B real-time usage, where the parameters of this model represent the various call-processing and overhead activities performed by the processor. Laboratory and test-site measurements of these parameters during development provided, through use of the real-time model, estimates of the call-processing capacity. In this way, any problem areas having an adverse effect on call-processing capacity could be identified as candidates for improvement during continued development. This same real-time model, at the completion of development, has been incorporated into the TSPSCAP program⁴ used by the operating telephone companies to determine the call-processing capacity of specific TSPS No. 1B sites. The formulation of this real-time model, the laboratory and test-site measurement techniques, and the resulting capacity performance data are described in subsequent sections of this paper.

1.2 System reliability analysis

An important step in the development of highly reliable switching systems is the prediction of their reliability. To provide uninterrupted service, TSPS No. 1B has the same reliability objectives as other Bell System electronic switching systems (ESSs), namely: an average downtime of less than 3.0 minutes per year.¹ In TSPS No. 1B, most of the TSPS peripherals are retained and their maintenance strategy remains virtually unchanged from the TSPS No. 1. Thus, the reliability objectives of the TSPS peripherals will not change in TSPS No. 1B from 1.0 minute per year average downtime and, consequently, the SPC 1B reliability must achieve the objective of less than 2.0 minutes per year average downtime.

To predict the reliability of SPC 1B hardware, continuous-time, finite-state Markov models were used. The Markov model approach for the reliability calculation of repairable systems is described in Ref. 5. Throughout the development period of TSPS No. 1B, the reliability model was updated to accurately reflect architectural modifications or design changes in the subsystems. The reliability estimates of various configurations were compared to monitor the system reliability, to identify limiting subsystems, and to determine if modifications would improve the overall reliability. This will be described in subsequent sections of this paper.

II. CAPACITY EVALUATION

2.1 Approach taken

The approach taken to model the real-time usage of the SPC 1B was to modify the real-time model of the SPC 1A.⁶ Emulation of the SPC 1A code by the SPC 1B made this approach possible. The modified model contains parameters that represent the speedup in SPC 1B instruction execution relative to the SPC 1A and, also, the effects of the DMERT operating system.

These modifications were characterized by making measurements of real-time usage at various call loads ranging from idle to over 160 percent of the SPC 1A capacity. Call loads were applied through use of electronic, programmable call generators attached to TSPS trunks. The response of TSPS operators to these calls was simulated by other electronic, programmable units. The measurements of real-time usage were made by non-interfering monitoring equipment, which sampled and recorded the system execution state every 10 microseconds. Other measurements consisted of various TSPS traffic counts periodically printed out on the standard output devices.

2.2 TSPS No. 1 capacity analysis

Because the SPC 1A real-time model forms the basis for the SPC 1B real-time model, it is briefly described here. References 6, 7, and 8 should be consulted for greater detail.

2.2.1 SPC 1A software architecture

During normal operation, most of the real-time usage of the SPC 1A occurs at two priority levels, called J-level and base level. J-level has the higher priority of the two, and is entered every 5 ms through a hardware interrupt to perform necessary input/output operations involved in communicating with the TSPS peripherals. Although a higher-priority H-level is also involved in these operations, H-level and J-level will hereafter be jointly referred to as J-level except where distinction is necessary. In the SPC 1A, base-level work has the lowest system priority and is performed whenever there are no higher-priority interrupts. Most of the call-processing work is performed in base level.

Each base-level program is assigned to one of five classes of work: A, B, C, D, or E. Each class is periodically visited by a control program to determine whether there is any work to do and to perform the work if present. The control program endlessly repeats the following fixed visitation sequence:

... ABACABADABACABABACABADABACABAE ...

We can see that from one class-E visitation to the next, termed an E-E cycle, the five classes are visited according to the ratio

$$A:B:C:D:E = 15:8:4:2:1.$$

Base-level programs are assigned to these classes in accordance with the acceptable delays in their execution; class A contains those programs requiring fastest response.

The time duration of an E-E cycle increases with the call load because, as the call load increases, there is more work to be done during each class visitation. However, a fixed amount of base-level work must be performed no matter what the call load is (e.g., determining if there is any work to do) and this work is referred to as the E-E cycle overhead.

2.2.2 SPC 1A real-time model

The real-time model developed for the SPC 1A consists of the equation

$$900 = t_N N + T_{CR} + t_E E. \quad (1)$$

This equation expresses how a quarter-hour (900 seconds) of processor real time is shared by three different kinds of processor work: trunk-seizure work, represented by $t_N N$; constant-rate work, represented by T_{CR} ; and E-E cycle overhead work, represented by $t_E E$. Each of these three terms is expressed in seconds per quarter-hour.

A TSPS call begins as a seizure (request for service) of a special TSPS trunk from a local office to a toll office. Most trunk seizures result in completed TSPS calls, but a few become uncompleted attempts because of customer abandonments, busy circuits, etc. Although these uncompleted attempts do not require as much processor real time as completed calls, they must be included as part of the processor real-time load. In the real-time equation, N represents the number of trunk seizures per quarter-hour; and t_N represents the average amount of processor real time (in seconds) required per trunk seizure. The value of t_N depends on the mix of various types of completed TSPS calls and uncompleted attempts. About two-thirds of t_N occurs in TSPS base level, and the other third in J-level. The TSPS call-processing capacity is expressed in terms of trunk seizures per quarter hour.

Constant-rate work is the processor work that is performed at fixed time intervals and is independent of trunk-seizure rate. For example, one type of TSPS trunk is scanned every 100 ms to determine whether a trunk seizure has occurred. The value of T_{CR} , in seconds per quarter hour, depends on the number of TSPS peripherals in use, and most of this time is spent in J-level.

The E-E cycle overhead work uses all processor real time not used by trunk seizures or constant-rate work. E represents the number of E-E cycles that are executed per quarter hour, and t_E represents the

average processor real time (in seconds) spent per E-E cycle in doing E-E cycle overhead work, which is independent of trunk-seizure load. By definition, all of this time occurs in base level.

Equation (1) is linear in terms of E and N . Figure 1 plots E as a function of N for a typical SPC 1A site. Such a plot is referred to as a load line, which describes how the E-E cycle rate, E , varies with respect to the trunk-seizure rate, N . The slope of this load line is $-t_N/t_E$ and the intercept, corresponding to an idle system (i.e., when $N = 0$), is $(900 - T_{CR})/t_E$.

Also shown in Fig. 1 is a value of E , called E_{MIN} , which is the lowest E-E rate that can be sustained while still providing adequate system response. At rates below E_{MIN} the visitation rate to the previously described base-level classes of work becomes too low and delays in serving requests become too long to meet service criteria. The trunk-seizure rate corresponding to E_{MIN} is defined as the quarter-hour trunk-seizure capacity, N_{CAP} .

2.2.3 TSPSCAP program

TSPSCAP is an interactive, time-shared program used by the operating companies to determine the trunk-seizure capacity of specific TSPS sites. The user inputs the call mix and hardware configuration of a site, and TSPSCAP calculates the values of t_N , T_{CR} , and E_{MIN} corresponding to these input values for use with the above real-time equation. TSPSCAP then outputs the value of N_{CAP} for that site, together with auxiliary information such as an equation for the E




Fig. 1—Typical SPC 1A load line.

versus N load line. This information is used by the operating telephone companies in growth planning to determine how close a TSPS No. 1 site is to its capacity limit.

2.3 TSPS No. 1B capacity analysis

As we mentioned earlier, emulation of the TSPS No. 1 software allows construction of the SPC 1B real-time model by modifying the SPC 1A real-time model. The modifications represent the speedup in instruction execution and the effects of the DMERT operating system. To understand these modifications the reader should know how the emulated code executes in the SPC 1B environment. This is briefly described below; Ref. 3 should be consulted for a more complete description.

2.3.1 SPC 1B software architecture

2.3.1.1 System execution levels. The DMERT operating system has sixteen execution levels (ELs), numbered 0 through 15, that determine the relative priorities for process execution; EL 15 has the highest priority. Kernel processes can use ELs 15 through 2, and supervisor/user processes are restricted to ELs 1 and 0. The emulated TSPS call-processing software executes as a kernel process.

Table I shows the ELs for those processes that influence the SPC 1B real-time usage. H-level and high-priority J-level of the emulated TSPS process execute at EL 12, and low-priority J-level executes at EL 11. Base level executes at EL 5. The DMERT timer, at EL 15, provides a timing function for other processes by notifying a requesting process after a specified time period has elapsed. Processes involved with I/O, file management, and memory management execute at ELs 10, 7, and 2, respectively. The scheduler at EL 2 schedules the supervisor/user processes at ELs 1 and 0. Diagnostics for the 3B20D and PSI execute at EL 0, whereas diagnostics for the TSPS peripherals remain as part of class-E work in emulated TSPS base level at EL 5. The new TSPS craft interface software, which uses DMERT facilities to provide maintenance input-output message capability and system

Table I—DMERT execution levels

Execution Level	Process
15	Timer
12, 11	TSPS J-Level
10	Disk Driver
7	File Manager
5	TSPS Base-Level
2	Memory Manager, Scheduler
1, 0	SPC 1B Diagnostics, TSPS Craft Interface, Other Supervisor/User Processes

status display, also executes at EL 0. Software for the other TSPS output messages (e.g., for periodic traffic counts) remains as part of the emulated J-level and base-level code.

A 3B20/DMERT timer hardware interrupt occurs every 10 ms to service any timing requests. As with the SPC 1A, a J-level interrupt occurs every 5 ms. These two interrupts are synchronized such that the J-level interrupts lead the timer interrupts by 1 ms. As with the SPC 1A, the emulated base level executes whenever nothing at a higher EL is executing, with the exception that base level periodically relinquishes control (goes to sleep) so as to allow processes at lower ELs to execute.

Just before going to sleep, base level requests that the DMERT timer wake it after a specified period has elapsed. The low-level processes at ELs 4 through 0 can then execute, subject to interrupts by processes at higher ELs. However, if all the low-level processes complete their work before the timer awakens base level, then base level is prematurely awakened by a software interrupt and the pending timer request is deactivated. Thus, any real time not needed by the low-level processes is given back to base level, which uses this real time to execute additional E-E cycles.

2.3.1.2 Speedup factors. The increased speed of the SPC 1B causes a net speedup in the execution of the emulated TSPS process relative to the SPC 1A. Not all portions of the emulated code experience the same degree of speedup, however, because of dependence on dynamic instruction mix, cache hit ratio, and ATB hit ratio.

The dependence on dynamic instruction mix occurs because some SPC 1A instructions could be emulated more efficiently than others. Also, in the SPC 1B, the execution time of some emulated SPC 1A instructions depends on what instruction options (e.g., rotating and masking) are exercised, whereas no such dependency exists in the SPC 1A.

To reduce memory access time, the SPC 1B employs a cache memory to contain the most recently accessed words of main memory. The cache is searched prior to each memory access and, if the word is in the cache (i.e., a cache hit), less real time is used because main memory need not be accessed. The cache is shared in common by all processes in the system.

The SPC 1B also employs eight Address Translation Buffers (ATBs), which speed up the task of translating from virtual memory address to physical memory address. Each ATB is essentially a cache memory that contains the physical addresses of the most recently accessed pages of virtual memory assigned to that ATB (a page is a 512-word block of main memory). If the page address is not in the ATB (i.e., an ATB miss), extra time is used in translation, which can

increase an instruction's execution time. To reduce ATB misses, J-level is exclusively assigned to one ATB and base level is exclusively assigned to another.

Parameters called "speedup factors" have been introduced to characterize the increased speed of the SPC 1B in executing the emulated TSPS process. Because of the effects of cache and ATB hits, speedup factors apply to execution of portions of code rather than to individual instructions. Thus, the speedup factor for a given portion of emulated code depends on the mix of executed instructions, and on the cache and ATB hit ratios experienced by those instructions.

2.3.1.3 DMERT operating system. For TSPS No. 1B, some real-time requirements of the DMERT operating system are application independent and others are application dependent. The application-independent requirements are for those functions that are necessary for maintaining a stable system environment. For example, the real time allocated to diagnose the 3B20D Processor would fall into this category. The application-dependent requirements are for those TSPS functions that make use of DMERT-supplied facilities. Two examples are: the real time required by the new TSPS craft interface, and the real time required to interface DMERT to TSPS J-level.

Parameters have been introduced that represent the combined TSPS-independent and TSPS-dependent DMERT real-time requirements for TSPS No. 1B. One parameter represents the combined high-level requirements (at ELs 15 through 5), and a second represents the combined low-level requirements (at ELs 4 through 0). Other parameters represent the real time used in handling TSPS J-level interrupts and base-level sleep requests.

2.3.2 SPC 1B real-time model

The SPC 1B real-time model is formed by adding speedup and operating system parameters to eq. (1) so as to obtain the new equation:

$$900 = t'_N N' + T'_{CR} + t'_E E' + T'_H, \quad (2)$$

where

N' = trunk-seizures per quarter hour serviced by the SPC 1B

E' = E-E cycles per quarter hour executed by the SPC 1B

T'_H = seconds per quarter hour used by high-level processes (at ELs 15 through 5) associated with TSPS-independent and TSPS-dependent DMERT work

and where t'_N , T'_{CR} , and t'_E are as defined in the following paragraphs.

The value of t'_N , the average processor seconds per trunk seizure, is defined as

$$t'_N = \frac{t_{NB}}{K_{NB}} + \frac{t_{NJ}}{K_{NJ}}, \quad (3)$$

where

t_{NB} = average processor seconds used per trunk seizure by the SPC 1A in base level

K_{NB} = SPC 1B speedup factor for t_{NB}

t_{NJ} = average processor seconds used per trunk seizure by the SPC 1A in J-level

K_{NJ} = SPC 1B speedup factor for t_{NJ} .

Separate base-level and J-level speedup factors are defined because base-level and J-level each has its own dynamic instruction mix and, also, its own ATB and associated ATB hit ratio.

The value of T'_{CR} , the processor seconds per quarter hour of constant-rate work, is defined as

$$T'_{CR} = \frac{T_{CRB}}{K_{CRB}} + \frac{T_{CRJ}}{K_{CRJ}} + 18 \times 10^4 t_{DJ}, \quad (4)$$

where

T_{CRB} = processor seconds per quarter hour used in constant-rate work by the SPC 1A in base level

K_{CRB} = SPC 1B speedup factor for T_{CRB}

T_{CRJ} = processor seconds per quarter hour used in constant-rate work by the SPC 1A in J-level

K_{CRJ} = SPC 1B speedup factor for T_{CRJ}

t_{DJ} = processor seconds used by the SPC 1B in handling each J-level interrupt.

Separate base-level and J-level speedup factors are defined for the same reasons stated above.

The value of t'_E , the processor seconds per E-E cycle to perform E-E cycle overhead work in base level, is defined as

$$t'_E = \frac{t_E}{K_E} + b(t_{DB} + \bar{s}A), \quad (5)$$

where

t_E = processor seconds used per E-E cycle by the SPC 1A in performing E-E cycle overhead work in base level

K_E = SPC 1B speedup factor for t_E

b = number of base-level sleep periods executed per E-E cycle by the SPC 1B

t_{DB} = processor seconds used by the SPC 1B in handling each base-level sleep-period request

\bar{s} = average duration (in seconds) of each base-level sleep period

A = average fraction of each base-level sleep period that is available to low-level processes (at ELs 4 through 0).

In this formulation it can be seen that the real time used by the low-level processes is treated as part of the SPC 1B E-E cycle overhead. The value of A decreases as high-level interrupts increase and, therefore, A decreases as call load increases.

The values of b and \bar{s} must satisfy the constraint

$$bs\bar{A}E^* = T'_L, \quad (6)$$

where

T'_L = seconds per quarter hour to be allocated to low-level processes (at ELs 4 through 0) associated with TSPS-independent and TSPS-dependent DMERT work

E^* = lowest SPC 1B E-E cycle rate at which T'_L is to be allocated by base-level sleep periods.

At E-E cycle rates less than E^* , insufficient base-level sleep periods will occur to satisfy eq. (6). At E-E cycle rates higher than E^* , more than T'_L can be used by low-level work if necessary.

The value of b is a software parameter, and the value of \bar{s} is determined by the value of s , which is another software parameter. When base level goes to sleep, it requests that it be awakened after s milliseconds have elapsed. Because this request can be made at any time relative to the 10-ms DMERT timer interrupt, \bar{s} is around 5 ms longer than s .

Equations (2) through (6) constitute the SPC 1B real-time model.

2.3.3 Determination of real-time model parameters

The newly introduced SPC 1B real-time parameters have been characterized through measurements made in the TSPS system laboratories and at the test site in Fresno, California, prior to cutover. The basic measurement technique involved measuring the percentage of processor real time used at each execution level under a number of different loads applied to the system. Other auxiliary measurements were also made.

2.3.3.1 Real-time measurement techniques. Processor real-time usage at the sixteen execution levels was measured through use of *Dynaprobe** monitoring equipment manufactured by the NCR COMTEN Corporation. The *Dynaprobe*, through means of high-impedance probes attached to the SPC 1B backplane, was used to sample the execution-level bits of the Program Status Word (PSW) every 10 microseconds to determine the relative frequencies of execution-level

* Registered trademark of NCR COMTEN Corporation.

occupancies. Other signals were also sampled at the same time to: (i) distinguish between emulated and native-mode code; (ii) count the number of times that each emulated SPC 1A instruction was executed during the measurement period; and (iii) measure the hit ratios experienced by the cache and by the base-level and J-level ATBs. The raw counts for all these data were written onto magnetic tape for subsequent off-line analysis.

During parameter measurement, simulated calls were generated by means of MICLOB (Microprocessor Controlled Load Box) units attached to the TSPS trunks. The response of TSPS operators, for those simulated calls that required operator assistance, was simulated by MOPS (Microprocessor Operator Position Simulator) units. The MICLOB and MOPS units are described in Ref. 9. Complete parameter characterization required taking measurements under various system conditions. Call loads were varied from zero to the maximum applicable simulated load. Different degrees of low-level activity were obtained by running processor and memory diagnostics and by causing different rates of output messages to be generated by the craft interface.

2.3.3.2 Measurement of speedup factors. Values for each of the defined speedup factors were calculated from measurements taken at the Fresno test site. Figure 2 shows the calculated values for each of the speedup factors plotted with respect to E' , the SPC 1B quarter-hour E-E cycle rate. The value of E' is inversely related to call load; $E' = 58,000$ corresponds to an idle system and $E' = 10,000$ corresponds to the maximum applied call load.

Figure 2 shows that the two J-level speedup factors, K_{NJ} and K_{CRJ} ,




Fig. 2—Speedup factors.

are constant with respect to E' . The value of K_{CRJ} is less than K_{NJ} because J-level constant-rate work makes heavier use of the instructions that have relatively low-emulation efficiencies. Although the cache and ATB hit ratios associated with both of these speedup factors were observed to slightly increase with decreasing E' (increasing call load), the effect of these changes was compensated for by a slight change in the dynamic instruction mix for K_{NJ} and, for K_{CRJ} , a higher percentage of conditional transfers taken.

For base level, the three speedup factors, K_E , K_{CRB} , and K_{NB} are seen to change with E' . The value of K_E , the speedup factor for E-E cycle overhead work, decreases with decreasing E' (increasing call load) because of a marked decrease in the base-level cache and ATB hit ratios as E' decreases. At zero call load, a relatively small portion of emulated code (the E-E cycle overhead work) is executed for a relatively high percentage of the time, causing the cache and ATB hit ratios to be at their highest values. The value of K_{CRB} , the speedup factor for base-level constant-rate work, was not measured directly but is set equal to K_E because this type of work is quite similar to E-E cycle overhead work and because only a small percentage of real time (less than 2 percent) is involved.

The value of K_{NB} , the speedup factor for base-level trunk-seizure work, is seen to increase with decreasing E' (increasing call load) even though the cache and ATB hit ratios are decreasing. This increase is caused by a decrease in the number of base-level instructions (excluding constant-rate and E-E overhead instructions) executed per trunk seizure as E' decreases. Figure 3 shows this effect. Measured values of I_{NB} , the number of base-level instructions executed per trunk seizure, are plotted versus E' . The dependence of I_{NB} is seen to be approximately linear with respect to E' over a wide range of values.

Investigation has indicated that this effect is at least partly caused by queueing for busy facilities (e.g., digit receivers). During each E-E cycle, if a queue exists, an attempt is made to remove all entries from the queue. Those entries that cannot be removed remain for the next E-E cycle, thereby causing extra instructions to be executed. As the call load increases, the probability of queue formation also increases. The E-E cycle rate decreases, however, thereby producing a net decrease in the number of base-level instructions executed per trunk seizure. This effect also occurs with the SPC 1A, but to a lesser degree because, as will be seen, the E-E cycle rate of the SPC 1A is lower than that of the SPC 1B when both are operating at the same trunk-seizure rate.

Curves were fitted to the calculated values of the speedup factors shown in Fig. 2 to obtain expressions for the parameters used in the SPC 1B real-time model. These expressions are:




Fig. 3—Base-level instructions for each trunk seizure versus E' .

$$K_{CRJ} = 1.70; K_{NJ} = 3.20$$

$$K_{CRB} = K_E = \frac{8.2762}{(2.7929 - 13.104 \times 10^{-6} E')}$$

$$K_{NB} = \frac{20.698}{(1 + 45.884 \times 10^{-6} E')(2.8649 - 14.179 \times 10^{-6} E')}$$

where, as previously defined, E' is the quarter-hour E-E cycle rate.

2.3.3.3 Measurement of DMERT real-time requirements for TSPS No. 1B. Parameters representing the real-time requirements of DMERT for TSPS No. 1B combine both TSPS-independent and TSPS-dependent work. TSPS-independent DMERT work includes DMERT functional work (e.g., audits, timer, etc.) and maintenance work associated with the SPC 1B (e.g., 3B20D diagnostics). TSPS-dependent DMERT work includes the TSPS craft interface work and work associated with handling the TSPS J-level interrupts and base-level sleep-period requests. These parameters were characterized by

Dynaprobe measurements of the real time used at each execution level and under various system operating conditions.

The TSPS craft interface real time is primarily used in producing output messages for maintenance purposes, and is a function of message rate, message length, and the number of output devices in use. The message rate is, in turn, a function of call load. Characterization involved measuring the real-time cost on a per-character basis and analyzing output messages generated by SPC 1A sites to determine representative message rates and lengths.

Measurements of T'_H , the parameter combining TSPS-independent and TSPS-dependent real-time requirements for high-level DMERT work, yielded

$$T'_H = 45.2 + 10^{-3} N' \text{ s/QH},$$

where N' is the TSPS No. 1B quarter-hour trunk-seizure rate. Measurements of T'_L , the parameter combining TSPS-independent and TSPS-dependent real-time requirements for low-level DMERT work, yielded

$$T'_L = 101.0 + 7.5 \times 10^{-3} N' \text{ s/QH}.$$

The value of T'_L is the amount of real time that should be allocated to achieve satisfactory execution of low-level activities under worst-case conditions (e.g., high maintenance activity during call overload). Under normal conditions, the actual value of T'_L is considerably less than this allocated value so that more real time is available to call processing.

To satisfy eq. (6), the values chosen for b , the number of base-level sleep periods per E-E cycle, and s , the requested duration of each sleep period, are

$$b = 5; s = 10 \text{ ms}.$$

Dynaprobe measurements also yielded

$$t_{DJ} = 84 \text{ microseconds}$$

for each J-level interrupt [see eq. (4)] and

$$t_{DB} = 1.5 \text{ ms}$$

for each base-level sleep request [see eq. (5)].

2.3.4 Model evaluation

Figure 4 shows measured and predicted values of the quarter-hour E-E cycle rate, E' , plotted versus the quarter-hour trunk-seizure rate, N' , for the Fresno TSPS site. The SPC 1B real-time model was used to predict three different E versus N load lines, each corresponding to a different low-level activity rate. The separate load lines occur be-




Fig. 4—SPC 1B load lines.

cause, as previously described, any real time not used by low-level processes is given back to base level, which uses this real time to execute additional E-E cycles.

The upper SPC 1B load line shows E versus N behavior when the low-level activity rate is low. Quarter-hour measurements were taken under these conditions at zero and the maximum applied load, and good agreement is seen between measured and predicted values. The middle SPC 1B load line corresponds to moderate low-level activity, and the lower SPC 1B load line corresponds to the condition when the low-level activity is heavy. Again, agreement between measured and predicted values is quite good.

2.3.5 SPC 1B capacity increase

Figure 5 shows two E versus N load lines that indicate the increase in call-processing capacity provided by the SPC 1B. The upper load line depicts the E versus N behavior for the Fresno TSPS site as predicted by use of the SPC 1B real-time model for a typical low-level activity rate. The lower load line shows the E versus N behavior of the




Fig. 5—Comparison between SPC 1A and SPC 1B load lines.

Fresno TSPS site as predicted using the SPC 1A real-time model, indicating how the site would perform if it were to use the SPC 1A. Both load lines assume the same call mix.

The SPC 1A load line shows that the SPC 1A would reach its capacity at about 5800 trunk seizures per quarter hour, since it is at that trunk-seizure rate that the E-E cycle rate equals 4180 E-Es per quarter hour, the SPC 1A value of E_{MIN} . Analysis and experiments conducted at Fresno indicate that E'_{MIN} for the SPC 1B should be less than E_{MIN} . Therefore, since SPC 1B measurements were conducted at Fresno at around 9300 trunk seizures per quarter hour with good system performance (see Fig. 4), it can be concluded that the capacity of the SPC 1B is at least 160 percent of the SPC 1A capacity. Furthermore, because the SPC 1B E-E cycle rate at 9300 trunk seizures per quarter hour is high with respect to the indicated value of E_{MIN} , it appears that the SPC 1B capacity is comfortably greater than 160 percent of the SPC 1A. This additional capacity serves as a margin to accommodate variation among sites with respect to call mix and peakedness in busy-hour load.

2.3.6 TSPSCAP program for TSPS No. 1B

A TSPSCAP program was developed for the TSPS No. 1B incorporating the SPC 1B real-time model. As has been seen, a load-line

equation for the SPC 1B is considerably more complex than for the SPC 1A, and depends to a large extent on the amount of SPC 1B diagnostic and craft interface activity. Therefore, instead of providing a load-line equation, the TSPSCAP program provides calculated values of E' and N' which can be used to plot two load lines for the site in question. These two load lines, similar to the upper and lower load lines shown in Fig. 4, define what can be termed as a load-zone of normal system behavior. That is, the quarter-hour E-E cycle and trunk-seizure measurements for a site that is experiencing normal operation should fall within this load zone.

III. RELIABILITY EVALUATION

3.1 Reliability requirements

The SPC 1B reliability requirements are similar to those of a traditional ESS-type processor, having four fault categories: hardware faults, recovery deficiencies, procedural errors, and software deficiencies.¹⁰

Hardware failures are allocated 0.4 minute of downtime per year. The SPC 1B is divided into three subsystems, each of which is duplicated to achieve high reliability. Thus, one failure in one side of a subsystem will not cause a system outage. Hardware faults can cause a system outage only when both sides of a subsystem are experiencing failures (i.e., before the first failure is repaired, another failure occurs on the other side of the subsystem). When this occurs, the system is unable to establish a working configuration until one side of the failed subsystem is repaired and system integrity is reestablished. The hardware reliability is a function of the failure rates of the subsystems, the system architecture, and the repair rates of the subsystems.

Recovery deficiencies are allocated 0.7 minute of downtime per year. When a hardware failure condition is detected, an automatic fault-recovery action occurs to establish a working configuration. Unsuccessful recovery actions are classified as recovery deficiencies. These are due to either design errors or limitations in fault-recovery programs.

Procedural errors are allocated 0.6 minute of downtime per year. An improper maintenance procedure can cause a system outage. Providing easy-to-follow documentation and reducing the number of manual steps help to minimize procedural errors.

Errors in operational programs and data are allocated 0.3 minute of downtime per year. The amount of bootstrap time required to recover the system from software deficiencies is considered to be a part of system downtime under this category. To minimize this source of downtime, overall software execution is monitored continually, data

integrity is checked using extensive auditing procedures, and thorough system integration tests are performed after program changes are introduced.

All four potential causes of system outage are closely interrelated. For example, improper procedures combined with certain hardware faults may prevent system recovery. In this paper, only SPC 1B outages induced by hardware faults are considered.

3.2 Reliability estimates

3.2.1 Reliability model

To provide a basis for the relationship between the reliability model and the system architecture, a brief review of the SPC 1B architecture is presented. A complete description of the SPC 1B architecture can be found in Ref. 11, and a more detailed description of 3B20D architecture can be obtained from Ref. 12.

As shown in Fig. 6, the SPC 1B consists of three subsystems or communities: a duplicated 3B20D Control Unit and PSI (CU/PSI), a duplicated Input/Output Processor (IOP), and a duplicated Disk File Controller with Movable Head Disk (DFC/MHD). Either half of the duplicated CU/PSI community can access either side of the duplicated TSPS peripheral bus system. The IOP community has duplicated

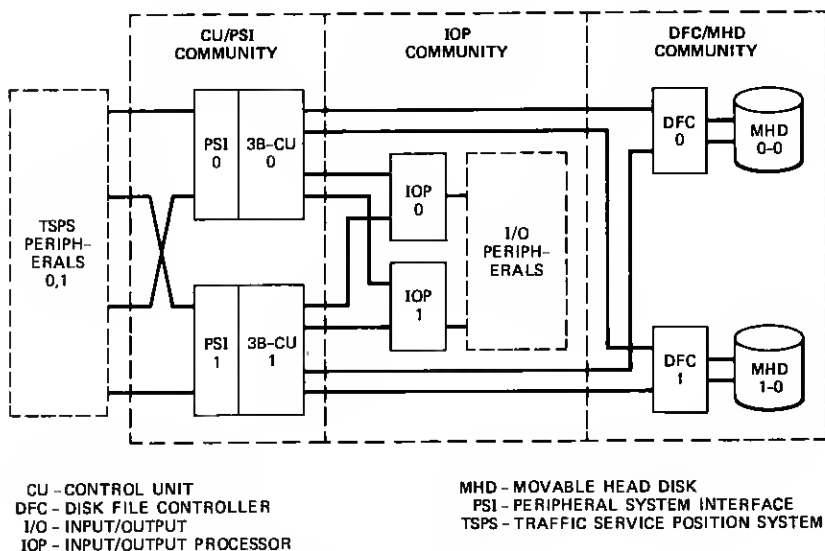


Fig. 6—SPC 1B architecture.

IOPs, each controlling a number of Peripheral Controllers (PCs). Each type of PC is designed to control a specific peripheral device, such as tape drives, teletypewriters, etc. The DFC/MHD community consists of a pair of DFCs, each capable of controlling one or more associated MHDs. One MHD on each half of the DFC/MHD community contains software required to bootstrap the SPC 1B.

Certain reliability measures are required to predict the probability of successful operation of the system. In expressing reliability of a switching system, "availability" is a more widely used term. Availability is defined as the fraction of time, on the average, that a system is expected to be in an operating condition. The availability of a switching system is a function of the system architecture and of subsystem failure rates and repair rates. Estimating availability of a system requires a mathematical model that can reflect the system architecture appropriately. The continuous-time, finite-state Markov model was used for the SPC 1B availability calculation, where an exponential probability distribution was assumed for the failure rates and repair rates. Detailed descriptions of how to use the Markov model to calculate system availability can be found in Refs. 5 and 13.

The reliability model for each of the three communities contains three states: a "duplex state," a "simplex state," and a "down state." The duplex state is a state where both halves of the community are fault-free and operational. Upon detecting a fault in either half of the community, a transition to the simplex state occurs. The rate of the transition is determined by the failure rate of the community. While a community is in the simplex state, one of two transitions is possible. A transition to the duplex state could occur if the faulty half of the community is repaired before a failure occurs in the other half of the community. On the other hand, a transition to the down state may occur if a new fault is detected in the remaining half before the initial fault is successfully repaired. A transition from the down state to the simplex state occurs when one of the faulty halves is repaired and put back to service. The rates of transitions, from the down state to the simplex state and from the simplex state to the duplex state, are determined by the repair rates, which are the reciprocals of corresponding mean time to repairs (MTTRs) for the community.

The probability that a community is in the down state is defined as the unavailability of the community. The SPC 1B is considered out of service when any of the CU/PSI, IOP, or DFC/MHD subsystems of the SPC 1B are in a down state. Hence, the unavailability of the system can be obtained by calculating the sum of unavailabilities of these three communities. The expected downtime per year for the SPC 1B can be estimated directly from the unavailability of the system.

3.2.2 Availability estimates and modifications

To evaluate the unavailability of each community, the reliability model was converted to a set of simultaneous equations where the unknowns are the probabilities of the states. Programs were written to solve the sets of equations corresponding to various architectural configurations. When these programs are used, sensitivity of the system downtime to the architectural variations as well as to the parametric values such as repair rates and failure rates of each community could be investigated. Coefficients of the equations were determined by the failure rates and repair rates of each community. Failure rates of the three communities were estimated from their component failure rates.

The repair rate of each community is estimated from the MTTR of mechanical failures, the MTTR of electrical failures, and a craft dispatch time. The MTTR of mechanical failures is considered separately from the MTTR of electrical failures because, for an MHD, the MTTR of mechanical failures is an order of magnitude longer than that of electrical failures. To minimize the MTTR, extensive diagnostic programs are included in the TSPS No. 1B, which can locate a fault within the resolution of three circuit packs. Detailed descriptions of diagnostic programs can be found in Refs. 11 and 14. A craft dispatch time is added to the MTTR when determining the repair rates of each community because the SPC 1B can be maintained by craft personnel located at a remote site. The dispatch time depends on the average travel time from the remote site and on the ratio between the average staffed hours and unstaffed hours per day of the TSPS No. 1B office.

The failure rates of the CU/PSI, DFC, and IOP are principally due to electrical failures. On the other hand, the failure rate of the MHD is due to roughly half mechanical and half electrical failures. Consequently, the MTTR of an MHD is much longer than the MTTRs of the other units. For the TSPS No. 1B application, the MTTR of an MHD has been improved through use of a spare MHD for each system.

Evaluation of the reliability model using current parameters shows that the reliability objectives for the TSPS No. 1B have been met.

IV. CONCLUSION

This paper has described the prediction and evaluation of the call-processing capacity and system reliability of the SPC 1B. The call-processing capacity has been estimated through means of a processor real-time model whose parameter values have been determined by laboratory and test-site measurements. The system reliability has been predicted through use of Markov modeling techniques. Performing this evaluation during TSPS No. 1B development to monitor progress was instrumental in meeting the capacity and reliability objectives.

V. ACKNOWLEDGMENTS

G. J. Kujawinski and the authors worked together as a team in evaluating the SPC 1B capacity; his contributions were essential to the success of the effort.

REFERENCES

1. R. E. Staehler and J. I. Cochrane, "Traffic Service Position System No. 1B: Overview and Objectives," B.S.T.J., this issue.
2. N. X. DeLessio and N. A. Martellotto, "Traffic Service Position System No. 1B: System Description," B.S.T.J., this issue.
3. R. J. Gill, G. J. Kujawinski, and E. H. Stredde, "Traffic Service Position System No. 1B: Real-Time Architecture Utilizing the DMERT Operating System," B.S.T.J., this issue.
4. L. B. Brisson and R. L. Potter, "A Time-Shared Program for Predicting TSPS No. 1 Capacity," Proc. Int. Conf. Commun., Seattle, 1973.
5. J. A. Buzacott, "Markov Approach to Finding Failure Times of Repairable Systems," IEEE Trans. Reliability, R-19 (November 1970), pp. 128-34.
6. R. J. Jaeger, Jr. and R. L. Potter, "Analysis of Processor Usage in Stored Program Controlled Telephone Switching Systems," Proc. Int. Conf. Commun., Philadelphia, PA, 1972.
7. Special issue entitled "Traffic Service Position System No. 1," B.S.T.J., 49, No. 10 (December 1970).
8. N. Farber, "A Model for Estimating the Real-Time Capacity of Certain Classes of Central Processors," Paper 426, Proc. Sixth Int. Teletraffic Congress.
9. J. J. Stanaway, Jr., J. J. Victor, and R. J. Welsch, "Software Development Tools," B.S.T.J., 58, No. 6, Part 1 (July 1979), pp. 1307-34.
10. P. W. Bowman, M. R. Dubman, F. M. Goetz, R. F. Kranzmann, E. H. Stredde, and R. J. Watters, "Maintenance Software," B.S.T.J., 56, No. 2 (February 1977), pp. 255-87.
11. G. T. Clark, H. A. Hilsinger, J. H. Tendick, and R. A. Weber, "Traffic Service Position System No. 1B: Hardware Configuration," B.S.T.J., this issue.
12. M. W. Rolund, J. T. Beckett, and D. A. Harms, "3B20D Processor & DMERT Operating System: 3B20D Central Processing Unit," B.S.T.J., 62, No. 1, Part 2 (January 1983), pp. 191-206.
13. M. D. Soneriu and D. S. Suk, "Markov Model for Estimating the Reliability of Duplicated and Repairable Computing Systems," Nineteenth Annual Tech. Symp., Pathways to System Integrity, Washington, D. C. (June 1980), pp. 87-95.
14. J. L. Quinn, F. M. Goetz, and R. L. Engram, "3B20D Processor & DMERT Operating System: Diagnostic Tests and Control Software," B.S.T.J., 62, No. 1, Part 2 (January 1983), pp. 367-81.

